Scheduling & Resource Management
Grid Forum Breakout Group Summary

1st Grid Forum Workshop

NASA Ames Research Center

June 16-18, 1999

Participants

Bill Nitzberg, NASA Ames / MRJ, bnitzberg@arc.nasa.gov
Mike Peterson, U. Florida, peterson@chem.ufl.edu
Jonathan Geisler, Northwestern U., geisler@ece.nwu.edu
Jon Weissman, U. Texas/U. Minnesota, jon@cs.utsa.edu
Jennifer Schopf, Northwestern U., jms@cs.nwu.edu
David Bader, U. New Mexico, dbader@eece.unm.edu
Greg Hommes, Lawrence Livermore Lab, hommes1@llnl.gov
Judy Beiriger, Sandia National Laboratories, jibeiri@sandia.gov
Barny Maccabe, U. New Mexico, maccabe@cs.unm.edu
Mitch Murphy, MHPCC, mitch@mhpcc.edu
Ed Hook, NASA Ames / MRJ, hook@nas.nasa.gov
Tom Cheatham, Harvard, cheatham@deas.harvard.edu
Dan Stefanescu, Harvard, dan@deas.harvad.edu
James P. Jones, NASA Ames / MRJ, jjones@nas.nasa.gov
Andy Yoo, LLNL, yoo2@llnl.gov
Cas Lesiak, NASA Ames / MRJ, clesiak@nas.nasa.gov
Bhroam Mann, NASA Ames / MRJ, bmann@nas.nasa.gov
Abdul Waheed, NASA Ames / MRJ, waheed@nas.nasa.gov
Mark Clement, BYU, clement@cs.byu.edu
Quinn Snell, BYU, snell@cs.byu.edu
Keith Jackson, LBNL, krjacks@lbl.gov
Gary Hoo, LBNL, gjhoo@lbl.gov
Hyo Jung Song, UCSD, hjsong@csag.ucsd.edu

Proposed Charter & Overall Goals

- Charter: “Solve grid resource management”
  - Co-chairs: Bill Nitzberg, Jennifer Schopf
- Goals
  - Better definitions of charter [by HPDC]
  - Progress in three areas identified
  - Work via mailing list: sched-wg@gridforum.org

Area: Advance Reservations

- Advance reservations
  - Capability: “reserve resources {R} for time period T”
  - Bill Nitzberg, chair
- Goals
  - Prototype reservation across different resource management packages (and sites) [SC 1999]
  - Specification for an API for advance reservations [Jun 2000]

Area: Super-scheduling

- Super-scheduling & “global queueing”
  - Capability: “given a job, run it on Grid resources”
  - Jenny Schopf, chair
- Goals
  - Prototype super-scheduler [SC 1999]

Area: Resource Specification & Semantics

- List of attributes/tokens (resource specification and semantics)
  - Language + tokens
  - Quinn Snell, chair
- Goals
  - list of attribute/value pairs [Oct 1999]
  - Specification for a common intermediate form for job description and resource specification [Jun 2000]

Working Notes
Resource Management
Breakout Group

Goals for this Breakout Group

- Goal for each "area" we want to focus on
- List of topics (maybe important), but we're not ready to attack them
  - e.g., naming

Levels of "Standardization"

- Architecture / models
  - e.g., picture of major parts and connections
- Capabilities / services / requirements
  - e.g., ability to make an advance reservation
- Languages / specification
  - e.g., resource description language
- Tokens / reserved words
- Interfaces / APIs
- Protocols

Resource Management

- “Queuing”
- Monitoring
- Scheduling
- Starting/stopping jobs (task management)
- Accounting and logging

Resource Management, cont.

- Resource & requirements specifications
  - both language & meaning of terms (e.g., “node”)
- Scheduling policy specification
- Status information

- Hooks to other facilities
  - authorization, allocations, accounting

What Could We "Talk About"? (Low-hanging Fruit)

- Advance reservations & co-scheduling
  - workshop held at Argonne last month (mostly focused on CPU)
- Resource specificiation
  - language & tokens
- Job scheduler API
  - esp. for "super" or "meta" schedulers
- Task management API
- Site requirements

Other Stuff to "Talk About"

- Resource fungability
- Description of scheduling policy
- Accounting records
- . . .

Architecture

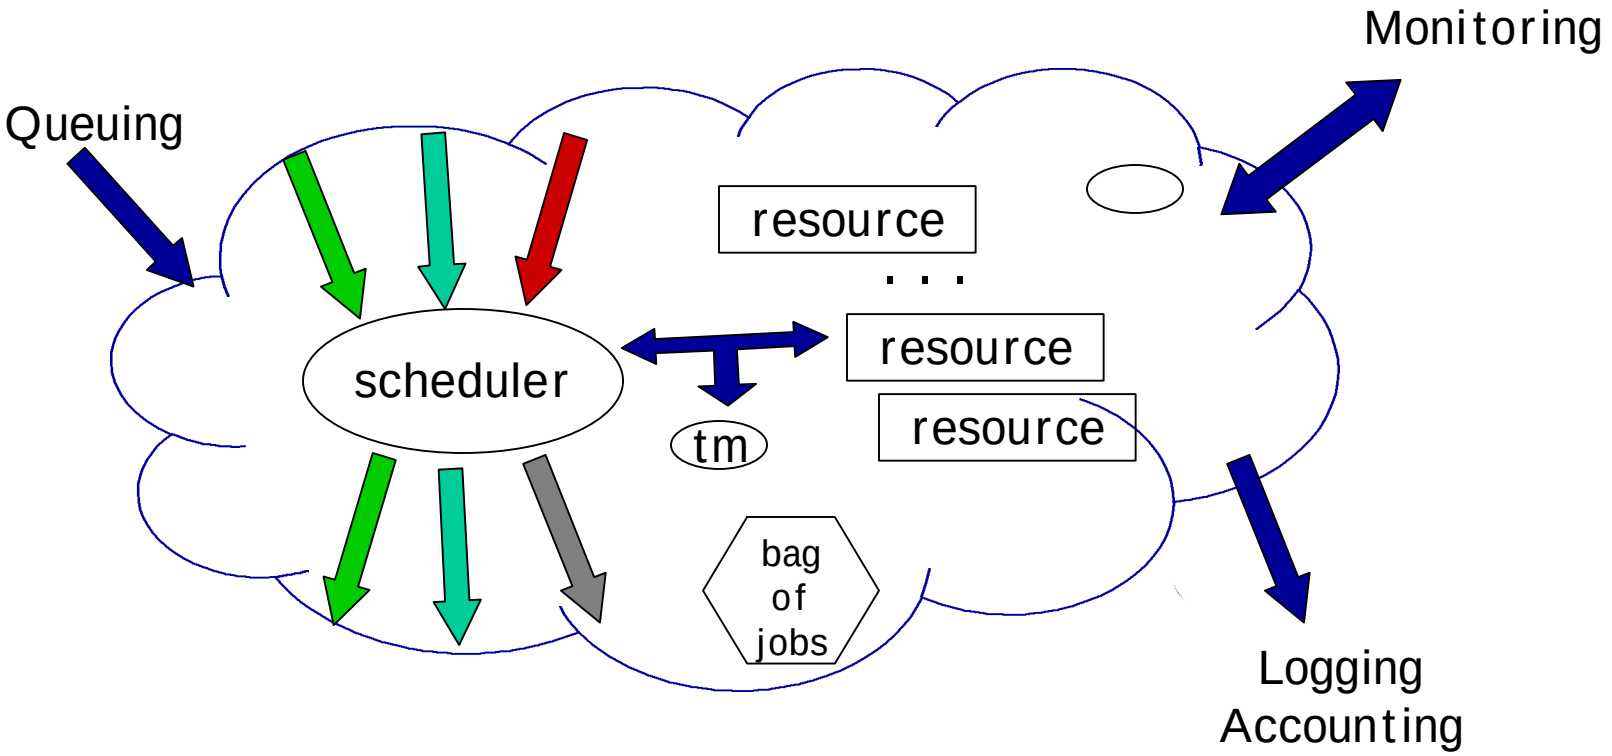

Scheduling Proposals

- Architecture

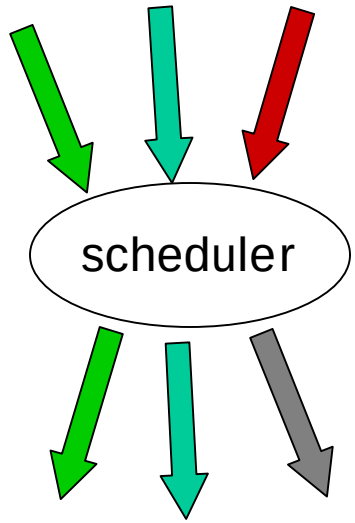

- Capabilities
  - advance reservations
  - start-time prediction
- Language
- Token
- Interfaces
  - to “meta” schedulers

Proposed “Scheduling” Goals

- List of capabilities of what we want schedulers to do
- List of capabilities schedulers need
  - e.g., specify dedicated access is required
  - ability to figure out which resources exist
  - ability to determine attributes/status of those resources

Proposed Scheduling Goals

- Taxonomy of scheduling approaches

Possible APIs

- super-scheduler to scheduler
- scheduler to peer-scheduler
- scheduler to resource manager
- scheduler to information sources
- scheduler to task manager
  - e.g, starting/stopping jobs, restarting
- monitoring interface into scheduler
- scheduler to application
  - dynamic resource allocation

Proposed Monitoring Goals

- Taxonomy of monitoring
  - e.g., job, system, user

Information Needed for Resource Management

- Specification language (and reserved words) for describing resources & describing what resources applications require
  - for schedulers, monitoring, qos
- Characterize different pieces of information w.r.t. semantics required
  - e.g., timeliness, precision

Interactions (?) with other Groups

- Accounting
- Information infrastructure
- Requirements for applications and tools

Other (important) things we didn't really talk about

- Accounting
- Queuing